



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/672,423	09/26/2003	John George Aschoff	IBMS.069PA(0510)	8426
63527 7590 05/22/2008				
DAVID W. LYNCH				
CHAMBLISS, BAHNER & STOPHEL				
1000 TALLAN SQUARE-S				
TWO UNION SQUARE				
CHATTANOOGA, TN 37402				
EXAMINER				
TSAL SIENG JEN				
ART UNIT		PAPER NUMBER		
2186				
MAIL DATE		DELIVERY MODE		
05/22/2008		PAPER		

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES

Ex parte JOHN GEORGE ASCHOFF, NEENAA CHERIAN,
and BRIAN JAY SMITH

Appeal 2008-0502
Application 10/672,423
Technology Center 2100

Decided: May 22, 2008

Before JOSEPH L. DIXON, JEAN R. HOMERE, and
ST. JOHN COURTENAY III, *Administrative Patent Judges*.

COURTENAY, *Administrative Patent Judge*.

DECISION ON APPEAL

This is a decision on appeal under 35 U.S.C. § 134(a) from the Examiner's rejection of claims 1-40. Claims 41-43 have been cancelled (App. Br. 2). We have jurisdiction under 35 U.S.C. § 6(b). We AFFIRM-IN-PART.

THE INVENTION

The disclosed invention relates generally to network storage systems. More particularly, Appellants' invention is directed to a method, apparatus and program storage device for providing automatic performance optimization of virtualized storage allocation within a network of storage elements (Spec. 1).

Independent claims 1 and 25 are illustrative:

1. An administration device for providing automatic performance optimization of virtualized storage allocation within a network of storage elements, comprising:

memory for storing data thereon; and

a processor configured for receiving from a user a request for storage of data, for determining workload requirements of the user making the request, for analyzing system parameters including performance characteristics of storage volumes within the network and for providing storage to meet the workload requirements of the user determined by the processor and to meet competing workload requirements based on the analysis of the system parameters.
25. A method for providing automatic performance optimization of virtualized storage allocation within a network of storage elements, comprising:

receiving from a user a request for storage of data;

determining workload requirements of the user making the request;

analyzing system parameters including performance characteristics of storage volumes within the network; and

providing storage to meet the determined workload requirements of the user and to meet competing workload requirements based on the analysis of the system parameters.

THE REFERENCES

The Examiner relies upon the following references as evidence in support of the rejections:

Lee	US 2003/0120864 A1	Jun. 26, 2003
Zahavi	US 6,820,035 B1	Nov. 16, 2004
		(filed Jun. 26, 2002)

THE REJECTIONS

Claims 1-9, 12-21, 24-33, and 36-40 stand rejected under 35 U.S.C. §102(e) as being anticipated by Zahavi.

Claims 10, 11, 22, 23, 34, and 35 stand rejected under 35 U.S.C. §103(a) as being unpatentable over Zahavi in view of Lee.

PRINCIPLES OF LAW

Appellants have the burden on appeal to the Board to demonstrate error in the Examiner's position. *See In re Kahn*, 441 F.3d 977, 985-86 (Fed. Cir. 2006).

In rejecting claims under 35 U.S.C. § 102, “[a] single prior art reference that discloses, either expressly or inherently, each limitation of a

claim invalidates that claim by anticipation.” *Perricone v. Medicis Pharm. Corp.*, 432 F.3d 1368, 1375-76 (Fed. Cir. 2005) (citation omitted).

ANALYSIS

Independent claims 1, 13, 39, and 40

We consider the Examiner’s rejection of independent claims 1, 13, 39, and 40 as being anticipated by Zahavi. Since Appellants’ arguments with respect to this rejection have treated these claims as a single group which stand or fall together, we select independent claim 1 as the representative claim for this rejection. *See* 37 C.F.R. § 41.37(c)(1)(vii).

Appellants contend that Zahavi does not disclose a processor configured for determining workload requirements of a user making a request for storage of data (App. Br. 7). In particular, Appellants contend that Zahavi discloses it is the *user* that determines the workload requirements rather than the *processor*, as claimed (App. Br. 8).

Appellants’ argument appears to narrowly interpret the language of the claim as being limited to a processor configured for determining the workload requirements of a user *responsive to the user* making a request for storage of data that is received by the processor. However, we note that the recited functional limitations of claim 1 do not require any particular sequence or order, as imputed by Appellants’ argument. Indeed, the “processor configured to” language of the claim does not positively recite the argued functional limitations. As claimed, the argued functional limitations may be performed at some future time, or not at all. Therefore, it is our reasoned view that the functional language argued by Appellants broadly reads on any ordering of functions, e.g., where the processor first

determines the workload requirements of a user, followed by the processor performing the function of receiving from the user a request for storage of data, or vice versa.

We find that Zahavi discloses the argued limitations of the processor receiving from a user a request for storage of data, as follows:

A write operation means data is requested to be placed on the disk whereas a read operation means data is requested to be viewed but not changed and typically this involves loading from disk or electronic memory such as cache.
(Zahavi, col. 7, ll. 16-20).

We find that Zahavi inherently discloses a processor and memory for storing user data as part of computer system 113 that enables the invention (Zahavi, col. 5, ll. 6-8) (“Software-based logic for enabling the invention resides on computer 113 (FIG. 2)”).

While we do not dispute Appellants’ observation that Zahavi discloses multiple examples where the user provides input to configure and manage the storage system (Reply Br. 2-3), we nevertheless find that Zahavi also discloses a processor configured for determining workload requirements of a user, as required by the language of representative claim 1.

It is our view that when the user presses the “Include” icon or button (Fig. 13, col. 9, ll. 60-62), the processor of Zahavi’s computer system then determines the workload requirements of the user that are reflected in the calculated results shown in summary table 536 (col. 9, ll. 60-62, *see e.g.*, the “Workload Index,” “Application ID,” “Data Type,” “Number of Devices,” and “IO Activity Adjusted for Protection” fields as shown in Fig. 13). We find these calculated metrics (Summary 536) broadly reflect the workload requirements of the user as a function of the configuration data previously

provided by the user in, e.g., entry fields 514 (“Performance Zone”), 520 (“Min. TB Needed”), and 516 (“Protection”), as shown in Fig. 13.

Moreover, our broad but reasonable reading of representative claim 1 on the Zahavi reference is fully consistent with Appellants’ Specification that discloses processing input/output in accordance with *a customer's specified performance and space requirements* (i.e., user input), as follows:

The administrator 270 includes a storage virtualization optimizer 272 that, according to an embodiment of the present invention, processes input/output *in accordance with a customer's specified performance and space requirements*, . . . [emphasis added].
(Spec. 10, ll. 14-17).

See also:

The virtualization optimizer 272 may make determinations of which nodes, i.e., engines such as the virtualization engine 274, may access the data, and which managed disk groups (groups of disks) would compose the LUNs to be selected. An additional important application of this would be to use the virtualization optimizer 272 to determine how to relocate, e.g., nodes or managed disk groups, the LUNs, i.e., virtual disks, *to meet the customer's desired level of performance* [emphasis added].
(Spec. 11, l. 22 – 12, l. 5).

See also:

The storage virtualization optimizer 272 allows *a user to point to a grouping of volumes and a particular window of time*, and then create a workload description based on the observed behavior of those volumes [emphasis added].
(Spec. 13, ll. 18-20).

While Appellants recite additional claim limitations on page 8 of the Brief, (i.e., “and to meet competing workload requirements based on the analysis of the system parameters), we note that Appellants do not offer any rationale or explanation as to why Zahavi purportedly does not disclose these additional limitations. Regarding these additional limitations, we conclude that Appellants have failed to traverse the Examiner’s rejection and instead have merely recited the language of the claim and asserted that such language is not taught by the reference. *See* 37 C.F.R. § 41.37(c)(1)(vii). Nevertheless, it is our view that Zahavi discloses meeting competing workload requirements (by using the processor to calculate the summary metrics discussed above) based on the analysis of the system parameters provided by the user via the user interface shown in Figure 13.

Based on the aforementioned reading of the claim on the reference, we find the weight of the evidence supports the Examiner’s finding that Zahavi is an anticipatory reference. On this record, we conclude that Appellants have not met the burden of coming forward with evidence or argument to show error in the Examiner’s *prima facie* case. Accordingly, we sustain the Examiner’s rejection of representative claim 1 (and independent claims 13, 39, and 40 that fall therewith) as being anticipated by Zahavi.

Independent claims 25 and 37

We consider next the Examiner’s rejection of independent claims 25 and 37 as being anticipated by Zahavi.

Since Appellants’ arguments with respect to this rejection have treated these claims as a single group which stand or fall together, we select

independent claim 25 as the representative claim for this rejection. *See* 37 C.F.R. § 41.37(c)(1)(vii).

In addition to the limitations we have previously addressed above for independent claim 1, Appellants further contend that Zahavi does not disclose providing automatic performance optimization of virtualized storage allocation within a network of storage elements by determining workload requirements of a user making a request for storage of data, as required by the language of claims 25 and 37 (App. Br. 9). Appellants further contend that Zahavi discloses a manual system, instead of an automatic system as claimed (*id.*).

We disagree. After reviewing the record before us, we conclude that Appellants have failed to specifically address the Examiner's findings on page 24 of the Answer that the aforementioned "providing automatic performance optimization" limitations are disclosed by Zahavi in the cited portions of columns 2, 6, and 7. While Appellants have cited other portions of the Zahavi reference in support of the proposition that Zahavi requires the user to identify workload requirements, Appellants have failed to respond to the Examiner's specific findings of fact.

We do not dispute Appellants' observation that Zahavi discloses multiple examples where the user provides (manual) input to configure and manage the storage system (*see* Reply Br. 2-3). Nevertheless, we find that Zahavi discloses automatic operation (i.e., calculations performed by the processor) that are initiated when the user presses the "Include" icon or button (Fig. 13, col. 9, ll. 60-62), as previously discussed.

For at least the aforementioned reasons, we find the weight of the evidence supports the Examiner's finding that Zahavi is an anticipatory

reference. On this record, we conclude that Appellants have not met the burden of coming forward with evidence or argument to show error in the Examiner's *prima facie* case. Accordingly, we sustain the Examiner's rejection of representative claim 25 (and independent claim 37 that falls therewith) as being anticipated by Zahavi.

Dependent claims 6 and 18

We consider next the Examiner's rejection of dependent claims 6 and 18 as being anticipated by Zahavi.

Appellants contend that Zahavi does not disclose that a processor determines workload attributes of the user and desired levels of performance, as claimed (App. Br. 10).

While we do not dispute Appellants' observation that Zahavi discloses multiple examples where the user provides (manual) input to configure and manage the storage system (*see* Reply Br. 2-3), we nevertheless find that Zahavi discloses a processor that determines workload attributes of the user and desired levels of performance (i.e., calculations performed by the processor) that are initiated when the user presses the "Include" icon or button (Fig. 13, col. 9, ll. 60-62), as previously discussed.

Because we conclude that Appellants have not met their burden of showing error in the Examiner's *prima facie* case, we sustain the Examiner's rejection of dependent claims 6 and 18 as being anticipated by Zahavi.

Dependent claims 7 and 19

We consider next the Examiner's rejection of dependent claims 7 and 19 as being anticipated by Zahavi.

For convenience, we reproduce the language of claim 7 here, noting that claim 7 and claim 19 recite equivalent limitations:

7. The administration device of claim 1, wherein the processor is configured for determining workload requirements of the user by using canned workload descriptions that are based on characterizations of user environments across various industries and applications.

Appellants contend that Zahavi does not disclose a processor that is configured for determining workload requirements of the user by using canned workload descriptions that are based on characterizations of user environments across various industries and applications, as required by the language of claims 7 and 19 (App. Br. 10). In contrast, Appellants assert that Zahavi discloses that the user determines workload requirements using knowledge of the application transactions or a Workload Library (*id.*).

The Examiner reads the claimed “canned workload descriptions” on Zahavi’s Library of Workloads 172 (Figure 5, column 8, line 42) (*see* Ans. 11; *see also* claims 7 and 19).

Zahavi discloses that the user provides information to the system (via the user interface shown in Figure 14) that relates to certain read/write characteristics, as follows:

Reference is now made to FIGS. 5 and 14. In using the invention, the user needs to provide through the user interface information identifying what percent of this work is Random-Read Hit, Random-Read Miss, Sequential Read and Write (discussed below). FIG. 14 shows an enlargement of screen area 524, including random Read Hits field 524a, Random Read Miss field 524b, Sequential Reads field 524c, And Writes field 524d. One approach to establishing these percentages is to first determine the Read/Write ratio as one way to establish

Read/Write Characteristics (FIG. 5, step 170). The user may determine this from knowledge of the application transactions or from a Workload Library (FIG. 5, step 172).

Once this ratio is established the user may attempt to determine what portion of the read activity is sequential. Sequential reads are generally almost 100% cache hits unless they come in bursts of small sequences. Again, this requires knowledge of the application or information from a Workload Library.
(Zahavi, col. 8, ll. 1-19).

We again do not dispute Appellants' observation that Zahavi discloses multiple examples where the user provides (manual) input to configure and manage the storage system (*see* Reply Br. 2-3). Nevertheless, it is our view that the language of claim 7 (and claim 19) broadly reads on Zahavi's disclosure of the user obtaining the relevant read/write characteristics or percentages (in one embodiment) from the Workload Library for entry into the user interface shown in Fig. 14, whereupon such entered data is used by the processor to determine the workload requirements of the user's various applications (i.e., "of the user"), as disclosed by Zahavi (col. 8, ll. 1-19). After a distinct I/O size is assigned to each type of I/O of the business unit, Zahavi further discloses the use of the "Include" button by the user to incorporate this new data into the summary table (col. 8, ll. 21-25). We find the use of the "Include" button invokes processor action, as previously discussed (*see also* Zahavi, Fig. 13, col. 9, ll. 60-62).

Because we conclude that Appellants have not met their burden of showing error in the Examiner's *prima facie* case, we sustain the Examiner's rejection of dependent claims 7 and 19 as being anticipated by Zahavi.

Dependent claims 8 and 20

We consider next the Examiner's rejection of dependent claims 8 and 20 as being anticipated by Zahavi.

Appellants contend that Zahavi does not disclose a processor that is configured for determining the workload requirements of the user by automatically creating workload requirements based on observations of storage access patterns of a user, as claimed (App. Br. 11).

After considering the record before us, we find the weight of the evidence supports the Appellants' position. While we have found *supra* that Zahavi discloses a processor that determines workload attributes (i.e., requirements) of the user and desired levels of performance that are calculated when the user presses the "Include" icon or button (Fig. 13, col. 9, ll. 60-62), we nevertheless acknowledge that Zahavi is silent regarding using *observations of storage access patterns of a user to automatically create workload requirements based on such observations*. To the contrary, as we have acknowledged above, Zahavi discloses multiple examples where the user provides (manual) input to configure and manage the storage system (*see* Reply Br. 2-3). Such manual user data is used as input data for the aforementioned "Include" icon or button function, as described above.

We note that the Workload Profiler 610 pointed to by the Examiner (Ans. 19-20) is an automated tool that correlates workload data on a storage system such that "affinity groups" of correlated logical devices may be defined for purposes of visualization (*see* Zahavi Fig. 29; *see also* col. 12, ll. 4-20). While such visual profiling of correlated logical devices may be suggestive of the claimed observation limitation, we nevertheless find that visualizing profiles of sets of correlated devices (Figs. 24-25) is not the same

as using a processor that is configured for determining workload requirements of the user by *automatically* creating workload requirements *based on observations of storage access patterns of a user*, as claimed. To the contrary, Zahavi discloses that the user (manually) references the aforementioned profiles to “set the amount of work to be used for workload clustering purposes.” (col. 14, ll. 17-18, step 834, Fig. 28). We note that “absence from the reference of any claimed element negates anticipation.” *Kloster Speedsteel AB v. Crucible, Inc.*, 793 F.2d 1565, 1571 (Fed. Cir. 1986). Because we conclude that Appellants have met their burden of showing error in the Examiner’s prima facie case, we reverse the Examiner’s rejection of dependent claims 8 and 20 as being anticipated by Zahavi.

Dependent claims 9 and 21

We consider the Examiner’s rejection of dependent claims 9 and 21 as being anticipated by Zahavi.

Dependent claims 9 and 21 recite a processor that is configured for determining workload requirements of the user by using intelligent software components that analyze workload descriptions for an application of the user.

Appellants contend that Zahavi merely mentions analyzing I/O activity, and thus does not disclose analyzing workload descriptions for an application of the user (App. Br. 12).

We disagree. We note that Zahavi’s invention is expressly directed to a “System and method for determining workload characteristics for one or more *applications* operating in a data storage environment” [emphasis added] (Zahavi, Title). Therefore, we find that Zahavi discloses a processor

that determines workload attributes of the user (i.e., of the user's applications) by using intelligent software components that analyze workload descriptions for an application of the user that are initiated when the user presses the "Include" icon or button (Fig. 13, col. 9, ll. 60-62), as previously discussed.

Because we conclude that Appellants have not met their burden of showing error in the Examiner's *prima facie* case, we sustain the Examiner's rejection of claims 9 and 21 as being anticipated by Zahavi.

Dependent claims 2-5, 12, 14-17, 24, 26-33, 36, and 38

We note that dependent claims 2-5, 12, 14-17, 24, 26-33, 36, and 38 were not separately argued in the Briefs. Therefore, we sustain the Examiner's rejection of these claims as being anticipated by Zahavi for the same reasons discussed *supra* with respect to representative independent claims 1, 13, and 25, respectively. *See* 37 C.F.R. § 41.37(c)(1)(vii).

Dependent claims 10, 11, 22, 23, 34, and 35

We note that dependent claims 10, 11, 22, 23, 34, and 35 were not separately argued in the Briefs. Therefore, we sustain the Examiner's rejection of these claims as being unpatentable over Zahavi in view of Lee for the same reasons discussed *supra* with respect to representative independent claims 1, 13, and 25, respectively. *See* 37 C.F.R. § 41.37(c)(1)(vii).

CONCLUSION OF LAW

Based on the findings of facts and analysis above, we conclude that Appellants have not met their burden of showing that the Examiner erred in rejecting claims 1-7, 9, 12-19, 21, 24-33, and 36-40 under 35 U.S.C.

§ 102(e) for anticipation. We also conclude that Appellants have not met their burden of showing that the Examiner erred in rejecting claims 10, 11, 22, 23, 34, and 35 under 35 U.S.C. § 103(a) for obviousness.

However, we conclude that Appellants have met their burden of showing that the Examiner erred in rejecting claims 8 and 20 under 35 U.S.C. § 102(e) for anticipation.

DECISION

We affirm the Examiner's decision rejecting claims 1-7, 9-19, and 21-40.

We reverse the Examiner's decision rejecting claims 8 and 20.

Appeal 2008-0502
Application 10/672,423

No time period for taking any subsequent action in connection with this appeal may be extended under 37 C.F.R. § 1.136(a)(1)(iv).

AFFIRMED-IN-PART

rwk

DAVID W. LYNCH
CHAMBLISS, BAHNER & STOPHEL
1000 TALLAN SQUARE-S
TWO UNION SQUARE
CHATTANOOGA TN 37402